

SERVICE SELECTION IN A SHARED ACCESS NETWORK SUPPORTING QUALITY OF SERVICE

John W. Garrett

Charles Robert Kalmanek Jr.

Han Q. Nguyen

Kadangode K. Ramakrishnan

Cross Reference to Related Applications

This application claims priority to United States Provisional Application Serial No. 60/190,633, entitled "INTERNET SERVICE SELECTION OVER CABLE," filed on March 20, 2000, and to United States Provisional Application Serial No. 60/190,636, entitled "QUALITY OF SERVICE OVER THE HFC CABLE PLANT," filed on March 20, 2000, the contents of which are incorporated by reference herein.

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60
61
62
63
64
65
66
67
68
69
70
71
72
73
74
75
76
77
78
79
80
81
82
83
84
85
86
87
88
89
90
91
92
93
94
95
96
97
98
99
100
101
102
103
104
105
106
107
108
109
110
111
112
113
114
115
116
117
118
119
120
121
122
123
124
125
126
127
128
129
130
131
132
133
134
135
136
137
138
139
140
141
142
143
144
145
146
147
148
149
150
151
152
153
154
155
156
157
158
159
160
161
162
163
164
165
166
167
168
169
170
171
172
173
174
175
176
177
178
179
180
181
182
183
184
185
186
187
188
189
190
191
192
193
194
195
196
197
198
199
200
201
202
203
204
205
206
207
208
209
210
211
212
213
214
215
216
217
218
219
220
221
222
223
224
225
226
227
228
229
230
231
232
233
234
235
236
237
238
239
240
241
242
243
244
245
246
247
248
249
250
251
252
253
254
255
256
257
258
259
260
261
262
263
264
265
266
267
268
269
270
271
272
273
274
275
276
277
278
279
280
281
282
283
284
285
286
287
288
289
290
291
292
293
294
295
296
297
298
299
300
301
302
303
304
305
306
307
308
309
310
311
312
313
314
315
316
317
318
319
320
321
322
323
324
325
326
327
328
329
330
331
332
333
334
335
336
337
338
339
340
341
342
343
344
345
346
347
348
349
350
351
352
353
354
355
356
357
358
359
360
361
362
363
364
365
366
367
368
369
370
371
372
373
374
375
376
377
378
379
380
381
382
383
384
385
386
387
388
389
390
391
392
393
394
395
396
397
398
399
400
401
402
403
404
405
406
407
408
409
410
411
412
413
414
415
416
417
418
419
420
421
422
423
424
425
426
427
428
429
430
431
432
433
434
435
436
437
438
439
440
441
442
443
444
445
446
447
448
449
450
451
452
453
454
455
456
457
458
459
460
461
462
463
464
465
466
467
468
469
470
471
472
473
474
475
476
477
478
479
480
481
482
483
484
485
486
487
488
489
490
491
492
493
494
495
496
497
498
499
500
501
502
503
504
505
506
507
508
509
510
511
512
513
514
515
516
517
518
519
520
521
522
523
524
525
526
527
528
529
530
531
532
533
534
535
536
537
538
539
540
541
542
543
544
545
546
547
548
549
550
551
552
553
554
555
556
557
558
559
560
561
562
563
564
565
566
567
568
569
570
571
572
573
574
575
576
577
578
579
580
581
582
583
584
585
586
587
588
589
590
591
592
593
594
595
596
597
598
599
600
601
602
603
604
605
606
607
608
609
610
611
612
613
614
615
616
617
618
619
620
621
622
623
624
625
626
627
628
629
630
631
632
633
634
635
636
637
638
639
640
641
642
643
644
645
646
647
648
649
650
651
652
653
654
655
656
657
658
659
660
661
662
663
664
665
666
667
668
669
670
671
672
673
674
675
676
677
678
679
680
681
682
683
684
685
686
687
688
689
690
691
692
693
694
695
696
697
698
699
700
701
702
703
704
705
706
707
708
709
710
711
712
713
714
715
716
717
718
719
720
721
722
723
724
725
726
727
728
729
730
731
732
733
734
735
736
737
738
739
740
741
742
743
744
745
746
747
748
749
750
751
752
753
754
755
756
757
758
759
760
761
762
763
764
765
766
767
768
769
770
771
772
773
774
775
776
777
778
779
780
781
782
783
784
785
786
787
788
789
790
791
792
793
794
795
796
797
798
799
800
801
802
803
804
805
806
807
808
809
810
811
812
813
814
815
816
817
818
819
820
821
822
823
824
825
826
827
828
829
830
831
832
833
834
835
836
837
838
839
840
841
842
843
844
845
846
847
848
849
850
851
852
853
854
855
856
857
858
859
860
861
862
863
864
865
866
867
868
869
870
871
872
873
874
875
876
877
878
879
880
881
882
883
884
885
886
887
888
889
890
891
892
893
894
895
896
897
898
899
900
901
902
903
904
905
906
907
908
909
910
911
912
913
914
915
916
917
918
919
920
921
922
923
924
925
926
927
928
929
930
931
932
933
934
935
936
937
938
939
940
941
942
943
944
945
946
947
948
949
950
951
952
953
954
955
956
957
958
959
960
961
962
963
964
965
966
967
968
969
970
971
972
973
974
975
976
977
978
979
980
981
982
983
984
985
986
987
988
989
990
991
992
993
994
995
996
997
998
999
1000
1001
1002
1003
1004
1005
1006
1007
1008
1009
1010
1011
1012
1013
1014
1015
1016
1017
1018
1019
1020
1021
1022
1023
1024
1025
1026
1027
1028
1029
1030
1031
1032
1033
1034
1035
1036
1037
1038
1039
1040
1041
1042
1043
1044
1045
1046
1047
1048
1049
1050
1051
1052
1053
1054
1055
1056
1057
1058
1059
1060
1061
1062
1063
1064
1065
1066
1067
1068
1069
1070
1071
1072
1073
1074
1075
1076
1077
1078
1079
1080
1081
1082
1083
1084
1085
1086
1087
1088
1089
1090
1091
1092
1093
1094
1095
1096
1097
1098
1099
1100
1101
1102
1103
1104
1105
1106
1107
1108
1109
1110
1111
1112
1113
1114
1115
1116
1117
1118
1119
1120
1121
1122
1123
1124
1125
1126
1127
1128
1129
1130
1131
1132
1133
1134
1135
1136
1137
1138
1139
1140
1141
1142
1143
1144
1145
1146
1147
1148
1149
1150
1151
1152
1153
1154
1155
1156
1157
1158
1159
1160
1161
1162
1163
1164
1165
1166
1167
1168
1169
1170
1171
1172
1173
1174
1175
1176
1177
1178
1179
1180
1181
1182
1183
1184
1185
1186
1187
1188
1189
1190
1191
1192
1193
1194
1195
1196
1197
1198
1199
1200
1201
1202
1203
1204
1205
1206
1207
1208
1209
1210
1211
1212
1213
1214
1215
1216
1217
1218
1219
1220
1221
1222
1223
1224
1225
1226
1227
1228
1229
1230
1231
1232
1233
1234
1235
1236
1237
1238
1239
1240
1241
1242
1243
1244
1245
1246
1247
1248
1249
1250
1251
1252
1253
1254
1255
1256
1257
1258
1259
1260
1261
1262
1263
1264
1265
1266
1267
1268
1269
1270
1271
1272
1273
1274
1275
1276
1277
1278
1279
1280
1281
1282
1283
1284
1285
1286
1287
1288
1289
1290
1291
1292
1293
1294
1295
1296
1297
1298
1299
1300
1301
1302
1303
1304
1305
1306
1307
1308
1309
1310
1311
1312
1313
1314
1315
1316
1317
1318
1319
1320
1321
1322
1323
1324
1325
1326
1327
1328
1329
1330
1331
1332
1333
1334
1335
1336
1337
1338
1339
1340
1341
1342
1343
1344
1345
1346
1347
1348
1349
1350
1351
1352
1353
1354
1355
1356
1357
1358
1359
1360
1361
1362
1363
1364
1365
1366
1367
1368
1369
1370
1371
1372
1373
1374
1375
1376
1377
1378
1379
1380
1381
1382
1383
1384
1385
1386
1387
1388
1389
1390
1391
1392
1393
1394
1395
1396
1397
1398
1399
1400
1401
1402
1403
1404
1405
1406
1407
1408
1409
1410
1411
1412
1413
1414
1415
1416
1417
1418
1419
1420
1421
1422
1423
1424
1425
1426
1427
1428
1429
1430
1431
1432
1433
1434
1435
1436
1437
1438
1439
1440
1441
1442
1443
1444
1445
1446
1447
1448
1449
1450
1451
1452
1453
1454
1455
1456
1457
1458
1459
1460
1461
1462
1463
1464
1465
1466
1467
1468
1469
1470
1471
1472
1473
1474
1475
1476
1477
1478
1479
1480
1481
1482
1483
1484
1485
1486
1487
1488
1489
1490
1491
1492
1493
1494
1495
1496
1497
1498
1499
1500
1501
1502
1503
1504
1505
1506
1507
1508
1509
1510
1511
1512
1513
1514
1515
1516
1517
1518
1519
1520
1521
1522
1523
1524
1525
1526
1527
1528
1529
1530
1531
1532
1533
1534
1535
1536
1537
1538
1539
1540
1541
1542
1543
1544
1545
1546
1547
1548
1549
1550
1551
1552
1553
1554
1555
1556
1557
1558
1559
1560
1561
1562
1563
1564
1565
1566
1567
1568
1569
1570
1571
1572
1573
1574
1575
1576
1577
1578
1579
1580
1581
1582
1583
1584
1585
1586
1587
1588
1589
1590
1591
1592
1593
1594
1595
1596
1597
1598
1599
1600
1601
1602
1603
1604
1605
1606
1607
1608
1609
1610
1611
1612
1613
1614
1615
1616
1617
1618
1619
1620
1621
1622
1623
1624
1625
1626
1627
1628
1629
1630
1631
1632
1633
1634
1635
1636
1637
1638
1639
1640
1641
1642
1643
1644
1645
1646
1647
1648
1649
1650
1651
1652
1653
1654
1655
1656
1657
1658
1659
1660
1661
1662
1663
1664
1665
1666
1667
1668
1669
1670
1671
1672
1673
1674
1675
1676
1677
1678
1679
1680
1681
1682
1683
1684
1685
1686
1687
1688
1689
1690
1691
1692
1693
1694
1695
1696
1697
1698
1699
1700
1701
1702
1703
1704
1705
1706
1707
1708
1709
1710
1711
1712
1713
1714
1715
1716
1717
1718
1719
1720
1721
1722
1723
1724
1725
1726
1727
1728
1729
1730
1731
1732
1733
1734
1735
1736
1737
1738
1739
1740
1741
1742
1743
1744
1745
1746
1747
1748
1749
1750
1751
1752
1753
1754
1755
1756
1757
1758
1759
1760
1761
1762
1763
1764
1765
1766
1767
1768
1769
1770
1771
1772
1773
1774
1775
1776
1777
1778
1779
1780
1781
1782
1783
1784
1785
1786
1787
1788
1789
1790
1791
1792
1793
1794
1795
1796
1797
1798
1799
1800
1801
1802
1803
1804
1805
1806
1807
1808
1809
1810
1811
1812
1813
1814
1815
1816
1817
1818
1819
1820
1821
1822
1823
1824
1825
1826
1827
1828
1829
1830
1831
1832
1833
1834
1835
1836
1837
1838
1839
1840
1841
1842
1843
1844
1845
1846
1847
1848
1849
1850
1851
1852
1853
1854
1855
1856
1857
1858
1859
1860
1861
1862
1863
1864
1865
1866
1867
1868
1869
1870
1871
1872
1873
1874
1875
1876
1877
1878
1879
1880
1881
1882
1883
1884
1885
1886
1887
1888
1889
1890
1891
1892
1893
1894
1895
1896
1897
1898
1899
1900
1901
1902
1903
1904
1905
1906
1907
1908
1909
1910
1911
1912
1913
1914
1915
1916
1917
1918
1919
1920
1921
1922
1923
1924
1925
1926
1927
1928
1929
1930
1931
1932
1933
1934
1935
1936
1937
1938
1939
1940
1941
1942
1943
1944
1945
1946
1947
1948
1949
1950
1951
1952
1953
1954
1955
1956
1957
1958
1959
1960
1961
1962
1963
1964
1965
1966
1967
1968
1969
1970
1971
1972
1973
1974
1975
1976
1977
1978
1979
1980
1981
1982
1983
1984
1985
1986
1987
1988
1989
1990
1991
1992
1993
1994
1995
1996
1997
1998
1999
2000
2001
2002
2003
2004
2005
2006
2007
2008
2009
2010
2011
2012
2013
2014
2015
2016
2017
2018
2019
2020
2021
2022
2023
2024
2025
2026
2027
2028
2029
2030
2031
2032
2033
2034
2035
2036
2037
2038
2039
2040
2041
2042
2043
2044
2045
2046
2047
2048
2049
2050
2051
2052
2053
2054
2055
2056
2057
2058
2059
2060
2061
2062
2063
2064
2065
2066
2067
2068
2069
2070
2071
2072
2073
2074
2075
2076
2077
2078
2079
2080
2081
2082
2083
2084
2085
2086
2087
2088
2089
2090
2091
2092
2093
2094
2095
2096
2097
2098
2099
2100
2101
2102
2103
2104
2105
2106
2107
2108
2109
2110
2111
2112
2113
2114
2115
2116
2117
2118
2119
2120
2121
2122
2123
2124
2125
2126
2127
2128
2129
2130
2131
2132
2133
2134
2135
2136
2137
2138
2139
2140
2141
2142
2143
2144
2145
2146
2147
2148
2149
2150
2151
2152
2153
2154
2155
2156
2157
2158
2159
2160
2161
2162
2163
2164
2165
2166
2167
2168
2169
2170
2171
2172
2173
2174
2175
2176
2177
2178
2179
2180
2181
2182
2183
2184
2185
2186
2187
2188
2189
2190
2191
2192
2193
2194
2195
2196
2197
2198
2199
2200
2201
2202
2203
2204
2205
220

SERVICE SELECTION IN A SHARED ACCESS NETWORK SUPPORTING QUALITY OF SERVICE

Field of the Invention

5 The present invention relates generally to communication network services, and, more particularly, to providing multiple services in a communication network.

Background of the Invention

10 Customers of communication network services often desire access to a plurality of different services and different service providers. For example, when using a dial-up connection to a packet-switched data network such as the Internet, a customer can choose from multiple service providers by dialing different telephone numbers in the PSTN. The physical path from the customer to
15 the customer's Internet Service Provider (ISP) is dedicated to the connection for the duration of the telephone call. The ISP assigns an IP address to the customer and can link the authenticated customer and the assigned IP address to the physical address (e.g. dial-up modem) used by the customer. With this linkage, the ISP can ensure the customer only uses the address authorized by the ISP and
20 can use the customer's IP address to manage access to the ISP's services. The physical connection between a customer and the ISP, as well as the linkage to IP address assignment and customer authentication is terminated when the dial-up connection is terminated.

 Constrained by the physical capacity of these temporary
25 connections across the PSTN, many service providers are moving to high-speed access architectures (e.g., digital subscriber line (DSL), wireless, satellite, or cable) that provide dedicated physical connectivity directly to the subscriber and under the control of the ISP. These alternatives to shared access through the switched telephone network, however, do not lend themselves to shared access by
30 multiple services and/or service providers.

Summary of the Invention

It is an object of the invention to enable multiple services or service providers to share the facilities of an access network infrastructure providing physical connectivity to subscribers. In accordance with an embodiment of the invention, a router situated at an edge of an access network forwards packets to any of a plurality of packet-switched service networks. The router uses a policy based on the source address of the packets to determine to which service network to forward the packet. Each network access device is assigned a network address, which is associated with a particular service or service provider to which the user of the device is subscribed. The network access device advantageously may be used in communication network services with a service or service provider that is separate from the operator of the access network infrastructure.

In accordance with another aspect of the invention, interconnections between a plurality of packet-switched service networks and an access network are localized into managed access points. Routers in the access network can advantageously forward packets to the managed access points using conventional routing procedures, thus enabling the access network to provide "local" packet-switched services. The managed access points use source address-based policy to determine to which service network to forward a packet. Where a packet arrives at a managed access point that is not connected to the correct service network, the managed access point can use packet encapsulation or some other form of tunneling to redirect the packet to the correct managed access point. The present invention, among other advantages, does not require interconnection points to each service network at every regional access network site.

In accordance with another aspect of the invention, a configuration server, upon receiving a request from a network access device selecting a particular service, allocates a network address from a pool of addresses associated with the service and assigns the network address to the network access device using a host configuration protocol, such as DHCP. In accordance with an embodiment of the invention, the configuration server authenticates the network

access device before assigning a network address. In accordance with another embodiment of the invention, the configuration server transmits authentication
65 information received from the network access device to a server in the service network responsible for authentication. In accordance with another embodiment of the invention, the host configuration protocol messages acknowledging allocation of a network address to the service subscriber's network access device are used to create entries in an address resolution protocol cache in order to
70 restrict access to the access network infrastructure to those network access devices that are properly registered and authenticated. In accordance with another embodiment of the invention, the host configuration protocol acknowledging allocation of the network address to the service subscriber's network access device are used to initiate the distribution of service policies to policy
75 enforcement points in order to enable differentiated quality-of-support for different subscribers of different services or service providers. The present invention, among other advantages, enables the network addresses—which ultimately determine the service network utilized by the particular network access device—to be allocated and reassigned dynamically.

80 These and other advantages of the invention will be apparent to those of ordinary skill in the art by reference to the following detailed description and the accompanying drawings.

Brief Description of the Drawings

85 FIG. 1 illustrates an interconnection of packet-switched service networks and an access network embodying principles of the invention.

 FIG. 2A and FIG. 2B is conceptual representation of an example embodiment illustrating principles of the invention based on an HFC access architecture with corresponding end-to-end protocol layers.

90 FIG. 3 is a flowchart of processing performed at a policy router, in accordance with an embodiment of the invention.

 FIG. 4 is a simplified example of router configuration instructions.

FIG. 5 illustrates an interconnection of packet-switched service network, regional access networks, and a packet-switched access network,
95 embodying principles of another aspect of the invention.

FIG. 6 is a flowchart of processing performed at a policy router acting as a managed access point, in accordance with another embodiment of the invention.

FIG. 7 is a timeline diagram of messages exchanged in the
100 assignment of a network address associated with a particular service to a network access device, in accordance with a preferred embodiment of another aspect of the invention.

FIG. 8 is a conceptual representation of a DHCP message exchanged between the network access device and a DHCP server.

FIG. 9 is timeline diagram of messages exchanged in the
105 assignment of a network address associated with a particular service to a network access device, in accordance with a preferred embodiment of another aspect of the invention.

FIG. 10 is a timeline diagram of messages exchanged in the
110 assignment of a network address associated with a particular service to a network access device, in accordance with a preferred embodiment of another aspect of the invention.

FIG. 11 is a flowchart of processing performed at a Cable Modem Termination System, exemplifying an embodiment of another aspect of the
115 invention.

FIG. 12 is a timeline diagram of messages exchanged in the assignment of a service class to a subscriber, in accordance with a preferred embodiment of another aspect of the invention.

FIG. 13 is a conceptual diagram of a hierarchical link-sharing
120 structure.

Detailed Description

In FIG. 1, a plurality of subscribers operating network access devices 101, 102, 103, ... 104 are provided access to communication network services, which are facilitated by a plurality of packet-switched data networks, shown in FIG. 1 as 151 and 152. Packet-switched data networks 151 and 152, referred to herein as "service networks," offer access to different services and/or are operated by different service providers. For example, service network 151 could provide packet-switched connectivity to public data networks while service network 152 could offer packet-switched telephony service (or the same public data network connectivity, but from a different service provider). The service networks, as is well known in the art, utilize a network addressing scheme to route datagrams to and from hosts: for example, where the service networks utilize the TCP/IP protocol suite, Internet Protocol (IP) addresses are assigned to each host and utilized in the process of routing packets from a source to a destination in the networks. See, e.g., "INTERNET PROTOCOL," IETF Network Working Group, RFC 791 (September 1981); S. Deering, R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification," IETF Network Working Group, RFC 1883 (December 1995), which are incorporated by reference herein. The invention shall be described herein with particular reference to the TCP/IP protocol suite and IP addresses, although those skilled in the art would readily be able to implement the invention using any of a number of different communication protocols.

The network access devices 101 ... 104 are typically customer premises equipment (CPE) such as a personal computer, information appliance, personal data assistant, data-enabled wireless handset, or any other type of device capable of accessing information through a packet-switched data network. Each network access device 101 ... 104 is either connected to or integrated with a network interface unit 111 ... 114, e.g. a modem, which enables communication through an access network infrastructure, shown as 120 in FIG. 1. Each network access device is assigned an IP address which, in accordance with an aspect of the invention, is associated with a particular service or service provider to which the user of the device is subscribed. For example, network access device 101 is

assumed to have been assigned, for purposes of the description herein, an IP address associated with a service provider operating service network 151. As
 155 further described herein, it is advantageous to provide a service activation system 160 which advantageously permits the dynamic allocation, assignment, and reassignment of IP addresses to the plurality of network access devices based on customer subscriptions to particular services.

The network access device 101 communicates with the service
 160 network 151 through the access network infrastructure 120, which, in accordance with aspects of the invention, is capable of recognizing and directing traffic to the proper service network. The access network infrastructure 120 advantageously can be operated and maintained by an entity that is the same as or different from the entities operating and maintaining the service networks 151 and 152. In
 165 accordance with an embodiment of an aspect of the present invention, the different IP-based services offered by the different service networks 151 and 152 utilize shared layer one and layer two resources in the access network 120. Layer three routing procedures, however, are modified to permit IP traffic from network access device 101 to flow to the correct subscribed service network 151. The
 170 access network 120 has a router 130 on the edge of the access network. The router 130 has a first interface with a connection to a router 141 in service network 151 and a second interface with a connection to a router 142 in service network 152. As further described herein, the router processes packets and is capable of directing traffic to the proper service network.

175 FIG. 2A shows an exemplary access architecture based on a hybrid fiber coaxial (HFC) access network. As is known in the art, each network interface device 201 ... 202 is either connected to or integrated with a cable modem 211 which enables communication through the HFC network 221. In accordance with the Data Over Cable Service Interface Specification (DOCSIS), a
 180 Cable Modem Termination System (CMTS), shown as 225 in FIG. 2A, communicates with the cable modems 211 and manages access to both upstream and downstream cable capacity on the HFC networks 221. See, e.g., "Data-Over-Cable Service Interface Specifications: Cable Modem Termination System –

Network Side Interface Specification,” Cable Television Laboratories, Inc., SP-
 185 CMTS-NSI-I01-960702; “Data-Over-Cable Service Interface Specifications:
 Cable Modem to Customer Premise Equipment Interface Specification,” Cable
 Television Laboratories, Inc., SP-CMCI-C02C-991015; “Data-Over-Cable
 Service Interface Specifications: Baseline Privacy Plus Interface Specifications,”
 Cable Television Laboratories, Inc., SP-BPI+-I06-001215, which are incorporated
 190 by reference herein. The CMTS 225 manages the scheduling of both upstream
 and downstream transmission and allocates cable capacity to individual customers
 identified by a Service IDs (SIDs). The CMTS 225 can have an integrated router
 228 or can be a separate device 226 that bridges to a fast Ethernet switch 227
 which connects to the router 228. The IP router 228 provides connectivity to an
 195 IP network 222, which further comprises the router 230 (corresponding to router
 130 in FIG. 1) which interfaces to IP routers 241 and 242 in service networks 251
 and 252, respectively. Accordingly, the HFC network 221, the CMTS 225, and
 the IP network 222 correspond to the access network infrastructure 120 shown in
 FIG. 1. FIG. 2B shows a conceptual diagram of the end-to-end communication
 200 protocol stack from a network access device 201 (101) to a router 241 (141) in
 service provider’s network 251 (151). As is known in the art, the lowest layer
 deals with the physical layer (PL) of the protocol stack, e.g. the Ethernet physical
 media device (PMD) layer; the second layer deals with the data link layer, e.g. the
 Ethernet Media Access Control (MAC) layer; and the third layer in the protocol
 205 stack deals with the network layer, e.g. the IP layer. The following aspects of the
 invention deal with modifications to routing processes in the network layer of the
 protocol stack.

Router 130 in the access network 120 in FIG. 1 (corresponding to
 IP router 230 in FIG. 2) separates the IP traffic to the multiple services or service
 210 providers as well as combines traffic from the multiple services or service
 providers. In accordance with an aspect of the invention, IP packets are routed
 from network access device 101 to the subscribed service network 151 using
 source address-based policy routing. Conventional routing is destination-based:
 the router consults an internal routing table which maps the destination addresses

215 of all inbound packets to a physical interface address for use for outgoing packets. Policy routing schemes, however, will selectively choose different paths for different packets even where the packet's destination address may be the same. Since network access devices are assigned addresses associated with a particular network service provider, the source address based policy routing scheme ensures
 220 packets from a network access device will go to the appropriate service network. Conventional destination-based routing will ensure that packets addressed to a network access device will be routed to the appropriate service network. Note that this would require service providers to advertise their service address ranges to their peers.

225 FIG. 3 sets forth the processing performed at a router in the access network, e.g. router 130 in FIG. 1. At step 301, the router receives an incoming packet. At step 302, the router reads the packet header and retrieves the packet filtering rules, typically stored in an access list as further described below. At steps 303, 305, and 307, the router applies the packet filtering rules. At step 303,
 230 the router compares the source IP address in the packet header to a list of addresses allocated to subscribers of services of a first service provider, e.g. operating service network 151 in FIG. 1. If the source address matches one of these addresses, then at step 304 the router forwards the packet to a router in service network 151, e.g. router 141 in FIG. 1. At step 305, the router compares
 235 the source IP address in the packet header to a list of addresses allocated to subscribers of services of a second service provider, e.g. operating service network 152 in FIG. 1. If the source IP address matches one of these addresses, then at step 305 the router forwards the packet to a router in service network 152, e.g. router 142 in FIG. 1. The router continues in this fashion with any other
 240 packet filtering rules identifying IP addresses allocated to subscribers of any other service providers. Assuming the IP source address does not match any such addresses associated with a service provider, at step 307, the router applies any remaining packet filtering rules and routes or denies the packet accordingly.

FIG. 4 sets forth an example of router configuration instructions
 245 written for the Cisco Internetworking Operating System (IOS), which is used

pervasively on conventional IP routers. Only the relevant portions of the configuration instructions are shown. Lines 401 to 405 configure the interface to utilize policy routing. Lines 406 to 410 specify the particular policy, namely to set the next “hop” address to the router address of a router in a one of the service networks, i.e. “isp1_next-hop_address,” if the source address matches a range of addresses allocated to subscribers of the services provided by the service network, i.e. “isp1_subs.” Lines 412 to 413 set forth access lists associating “isp1_subs” with ranges of addresses expressed, by convention, as a source address and a mask portion, i.e., the above policy is applied by the router to any traffic with a subscriber source address expressed as “isp1_prefix1” with a mask portion of “isp1_prefix1_wildcard”.

The embodiment shown in FIG. 1 notably requires interconnection points to all relevant service networks at each edge of each regional access network. In accordance with another aspect of the invention, it is desirable to create a regional transport network of routers and to localize the interconnection between the service networks and the access infrastructure into managed access points. A managed access point is a physical location at which the interfaces to the service networks can be provided. Having one or a small plurality of managed access points advantageously allows service selection to be implemented without requiring network service providers to connect physical facilities into, for example, every cable head end in an HFC-based network—thereby reducing costs for both the access network infrastructure operator and the service network providers. Each router in the regional transport network can be configured with policy information and invoke source address routing to forward packets to the managed access point providing access to the relevant service network. By overriding normal routing procedures, however, these procedures may introduce potential routing loops absent significant coordination between the routers external to known routing protocols. This risk can be minimized by centralizing the policy routing function in a single router that provides the interfaces to the service networks.

FIG. 5 illustrates an embodiment of this aspect of the invention. Each network access device 501 is connected through a network interface unit 511 to one of a plurality of access networks, e.g. 521 and 522 in FIG. 5. Each access network has an edge router (531 and 532 respectively in FIG. 5) which connects the access network to a regional IP network of routers, represented abstractly in FIG. 5 as IP access network 570. It is advantageous to aggregate connections from groups of edge routers to a single aggregation router 571 in the IP access network 570, as shown in FIG. 5. Aggregation router 571 can then connect to other routers in the regional IP network 570, i.e. routers 572 ... 573, which can also be aggregation routers connecting to pluralities of edge routers. Routers 541, 542, ... 543 in service networks 551, 552, ... 553 connect to the IP network 570 at routers 574 and 575, which act as managed access points to the service networks. Only the managed access point routers, e.g. 575, need invoke policy routing based on packet source address. All intermediate routers within the IP access network 570, i.e. 571 ... 573, use normal destination-based forwarding procedures for destinations that are not local to the network 570. No configuration of policy in the intermediate routers is necessary.

By locating the policy routing functions at the interfaces to the service networks, the access network infrastructure (whether reflected generally by 120 in FIG. 1 or, in the packet-switched context, as network 570 in FIG. 5) can provide access to "local" services available from within the access network infrastructure. For example, IP network 570 can provide access to "local" packet-switched services and operate independent of the source address assigned to the network access devices. Since the intermediate routers 571 ... 573 all use conventional destination-based forwarding, network 570 will properly route local service packets along the correct routing paths. "Non-local" service packets, however, are routed towards the managed access point routers 575 and 574 and policy routed to the correct service network. Where the network 570 forwards to a single managed access point router or where each managed access point router has a connection to each service network, the managed access point router can forward packets in accordance with the policies described above. Where,

however, there are service networks that connect to only a subset of the managed access points (e.g., in FIG. 5, where service network 553 only connects to managed access point router 574), packets can be redirected or “tunneled” to the correct managed access point in order to ensure that the packets arrive at the correct service network. The multiple interconnected managed access points can then provide a single “logical” inter-domain gateway, again permitting all other routers to use conventional destination-based routing procedures.

FIG. 6 sets forth a flowchart of the processing performed at a managed access point router, e.g. router 575 in FIG. 5, illustrating an embodiment of this aspect of the invention. The particular managed access point router 575 is assumed to be connected to two service networks, e.g. service networks 551 and 552 in FIG. 5, while a second managed access point router 574 provides access to a third service network, service network 553. At step 601, the router receives an incoming packet. At step 602, the router reads the packet header and retrieves the packet filtering rules, as well as decapsulates any encapsulated packets, as further described herein. At steps 603, 605, 607, and 609, the router applies the packet filtering rules. At step 603, the router compares the source IP address in the packet header to a list of addresses allocated to subscribers of services of a first service provider, e.g. operating service network 551 in FIG. 5. If the source address matches one of these addresses, then at step 604 the router forwards the packet to a router in service network 551, e.g. router 541 in FIG. 5. At step 605, the router compares the source IP address in the packet header to a list of addresses allocated to subscribers of services of a second service provider, e.g. operating service network 552 in FIG. 5. If the source IP address matches one of these addresses, then at step 606 the router forwards the packet to a router in service network 552, e.g. router 542 in FIG. 5. At step 607, the router compares the source IP address in the packet header to a list of addresses allocated to subscribers of services of a third service provider, e.g. operating service network 553 in FIG. 5, which is not connected to this particular managed access point. If the source IP address matches one of these addresses, then at step 608, the router encapsulates the packet, using any of a number of known methods for packet

encapsulation, and routes the packet to a new destination address, namely the address of the managed access point with access to service network 553, i.e. managed access point router 574. Packet encapsulation is a method by which a packet may rerouted to an intermediate destination other than the destination that would be selected using normal routing procedures. See, e.g., C. Perkins, "IP Encapsulation within IP," IETF Network Working Group, RFC 2003 (October 1996); C. Perkins, "Minimal Encapsulation within IP," IETF Network Working Group, RFC 2004 (October 1996), which are incorporated by reference herein. The receiving router 574 will decapsulate the packet and route the packet, accordingly, to service network 553. The router continues in this fashion with any other packet filtering rules identifying IP addresses associated with any other service providers. Assuming the IP source address does not match any addresses associated with any other service providers, at step 609, the router applies any remaining packet filtering rules and routes or denies the packet accordingly. Note that if managed access point 575 has a direct physical connection to managed access point 574, then no encapsulation is needed. In fact, a typical configuration might include multiple port-constrained policy routers on a GIG Ethernet providing the logical managed access point function without any encapsulation. Encapsulation is only really needed to provide a logical direct connection if there is not a direct physical connection.

Packets traveling between network access devices connected to the same access network infrastructure can be forwarded directly between the devices in the access network – rather than forwarding the packets outwards to a service network and back to the same access network. This advantageously saves on bandwidth and other network resources. The only packets that need be routed to a managed access point router need be the ones for which no specific route is known internally to the access network infrastructure.

It is advantageous to enable the IP addresses—which ultimately determine the service network utilized by the particular network access device—to be allocated and reassigned dynamically. With reference to FIG. 1, a service activation system 160 is shown which further comprises a configuration server

161 and a registration server 162 connected to the access network infrastructure
 370 120. The registration server 162 provides a network-based
 subscription/authorization process for the various services shared on the access
 network infrastructure 120. A customer desiring to subscribe to a new service can
 access and provide registration information to the registration server 162, e.g. by
 using HTML forms and the Hyper Text Transfer Protocol (HTTP) as is known in
 375 the art. Upon successful service subscription, the registration server 162 updates a
 customer registration database 163. The configuration server 161 uses the
 registration information to activate the service. The configuration server 161 is
 responsible for allocating network addresses on behalf of the service networks
 from a network address space associated with the selected service. In a preferred
 380 embodiment of this aspect of the invention, the configuration server 161 uses a
 host configuration protocol such as the Dynamic Host Configuration Protocol
 (DHCP) to configure the network addresses of the network access devices. See
 R. Droms, "Dynamic Host Configuration Protocol," IETF Network Working
 Group, RFC 2131 (March 1997); S. Alexander, R. Droms, "DHCP Options and
 385 BOOTP Vendor Extensions," IETF Network Working Group, RFC 2132 (March
 1997); which are incorporated by reference herein. This aspect of the invention
 shall be described herein with particular reference to DHCP, and the configuration
 server 161 shall be referred to herein as the DHCP server, although those skilled
 in the art would readily be able to implement this aspect of the invention using a
 390 different protocol.

FIG. 7 is a simplified timeline diagram of DHCP messages
 exchanged as the DHCP server 720 assigns a service-specific network address to a
 network access device 710 acting as a DHCP client. At 701, the network access
 device 710 sends a DHCPDISCOVER message through the access network
 395 infrastructure. The DHCPDISCOVER message, in accordance with an aspect of
 the invention, includes a "svc-id" option field that identifies the service to which
 the network access device has been subscribed and from which service is desired.
 The DHCP server 720 receives the DHCPDISCOVER message and, at 702,
 allocates an IP address from the pool of addresses associated with the particular

400 service. The DHCP server 720 can use the device's MAC address to lookup the customer's registration information and confirm that the device is authorized to access the identified service. Where the DHCP server 720 cannot find the device's MAC address in the registration database, the server can allocate a special IP address that only allows access to the registration server. At 703, the
 405 DHCP server 720 responds with a DHCPOFFER message that includes the IP address in a field in the DHCP message. At 704, the network access device 710 receives the DHCPOFFER (and any other offers from any other DHCP servers in the access network) and sends out a DHCPREQUEST directed to the DHCP server which requests the IP address identified in the DHCPOFFER. At 707, the
 410 DHCP server commits to assigning the IP address to the network access device, commits the binding to persistent storage, and transmits a DHCPACK message containing the configuration parameters for the device. If the DHCP server is unable to satisfy the DHCPREQUEST message, the server responds with a DHCPNAK message. FIG. 8 is a simplified representation of the fields in a
 415 DHCP message, including a "svc-id" option field (820) which identifies the selected service.

It is preferable that the DHCP servers and clients use some mutual authentication mechanism to restrict address assignment to authorized hosts and to prevent clients from accepting addresses from invalid DHCP servers. See, for
 420 example, the "delayed authentication" scheme described in R. Droms, W. Arbaugh, "Authentication for DHCP Messages," IETF Network Working Group, Internet Draft, <draft-ietf-dhc-authentication-____.txt>; or the Kerberos-based authentication mechanism described in K. Hornstein, T. Lemon, B. Aboba, J. Trostle, "DHCP Authentication via Kerberos V," IETF Network Working Group,
 425 Internet Draft, <draft-hornstein-dhc-kerbauth-____>; which are incorporated by reference herein. The "delayed authentication" mechanism supports mutual authentication of DHCP clients and servers based on a shared secret, which may be provisioned using out-of-band mechanisms. On the other hand, the Kerberos-based mechanisms are very well suited for inter-realm authentication, thereby
 430 supporting client mobility, i.e. a network access device could connect to a

particular access network infrastructure without any prior registration with the access network. Each service network provider could securely authenticate the network access device accessing the service network from another network “realm,” e.g. the access network infrastructure.

435 The operator of the relevant service network, e.g. service network 152 in FIG. 1, may desire to maintain a separate registration server, e.g. server 155 in FIG. 1, and to retain responsibility for user authentication and authorization. The service activation system 160 can provide a proxy server configured to permits HTTP traffic only between local hosts and registration

440 server 155 in service network 152. The service provider operating service network 152 would then be responsible for providing the appropriate registration information required for proper service selection to the service activation system 160. The service provider would also be responsible for notifying the service activation system 160 when service should be discontinued. Alternatively, the

445 DHCP server 161 in the service activation system 160 can interact with the registration server 155 using a back-end authentication protocol, e.g. the Remote Authentication Dial In User Service (RADIUS). See C. Rigney, A. Rubens, W. Simpson, S. Willens, “Remote Authentication Dial In User Service (RADIUS),” IETF Network Working Group, RFC 2058 (January 1997), which is incorporated

450 by reference herein. The DHCP server can contain a RADIUS client and, thereby, leverage the large RADIUS embedded base used for dial access authentication. FIG. 9 illustrates this embodiment of this aspect of the invention in a flowchart corresponding to the flowchart shown in FIG. 7. At 903, the DHCP server 920 generates a random challenge and includes the challenge along with the allocated

455 IP address in the DHCPOFFER message. The DHCP client 910 generates a response to the challenge by encrypting the challenge with a key that is derived from the subscriber’s authentication information. At 904, the client 910 includes the challenge, response, and IP address in the DHCPREQUEST message. The DHCP server 920 forwards both the challenge and response in a

460 RADIUS_ACCESS_REQ message to a RADIUS server 930 in the selected service network. The RADIUS server 930 either accepts or rejects the RADIUS

request and responds accordingly at 906. If the RADIUS request is accepted, the DHCP server 920 sends a DHCPACK message at 907 and the client 910 enters a bound state. If the RADIUS request is rejected, the DHCP server 920 sends a
 465 DHCPNACK message which informs the client 910 that the IP address that was allocated has been withdrawn.

Rather than modifying the DHCP protocol in order to establish the association between the device's MAC address and the service selection, it is advantageous to use an alternative two-step process. The registration server 162
 470 can associate the customer's IP address (used for registration) with the subscribed service, while the configuration server 161 can associate the customer's IP address with the customer's MAC address. The servers can interact and share their associations in the registration database 163 shown in FIG. 1. Thus, the subscriber can register the service selection with the registration server which
 475 temporarily establishes the association between the network access device's hardware address (e.g. the MAC address of the device) and the chosen service selection. The configuration server then uses the MAC address of the network access device to assign an IP address from the proper address space. FIG. 10 is a simplified timeline diagram of DHCP messages exchanged, in accordance with
 480 such an embodiment. At 1001, the network access device 1010 registers a service selection with the registration server 1030. It is assumed that the subscriber has passed the proper authentication procedures for the particular service selected, either beforehand (e.g. through transactions directly with the service provider's network) or in the same session with the registration server. The registration
 485 server 1030 sends some acknowledgment 1002 to the network access device 1010. At 1003, the registration server 1030 stores the selected service and associates the service selection with the hardware device address (e.g. MAC address) of the network access device 1010. After receiving the acknowledgment from the registration server 1030, the network access device 1010 releases any pre-existing
 490 address assignment by issuing a DHCPRELEASE message at 1004. At 1005, the network access device issues a standard DHCPDISCOVER message (i.e., there is no need for the service selection id in the message set forth above). The DHCP

server 1020 receives the DHCPDISCOVER message and, at 1006, allocates an IP address from the pool of address associated with the particular service associated
 495 with the device's hardware address. At 1007, the DHCP server 1020 sends a DHCPOFFER message that includes the IP address in a field in the DHCP message. At 1008, the network access device 1010 receives the DHCPOFFER and sends out a DHCPREQUEST back to the DHCP server 1030. At 1009, the DHCP server 1030 commits to assigning the IP address to the network access
 500 device 1010, commits the binding to persistent storage, and transmits a DHCPACK message containing the configuration parameters for the device.

It is desirable to restrict access to the access network infrastructure to those network access devices that are properly registered and authenticated. In accordance with another aspect of the invention, the access network infrastructure
 505 can be configured to perform access control taking advantage of the above-described address allocation process. An access network infrastructure with broadcast capabilities will often use a protocol such as the Address Resolution Protocol (ARP) to map network layer addresses used by the packet-switched networks to the hardware addresses used in the datalink layer of the access
 510 network infrastructure. See, e.g., D. Plummer, "An Ethernet Address Resolution Protocol," IETF Network Working Group, RFC 826 (November 1982). For example, and with reference to the HFC embodiment shown in FIG. 2, the CMTS 225 has an ARP cache, which is a table of entries storing bindings between IP addresses and the hardware MAC addresses assigned to network access devices
 515 201, 202. The ARP cache permits the CMTS to learn the correspondence between IP addresses and MAC addresses without resorting to broadcasting an ARP request. As is known in the art, the CMTS can act as a DHCP relay agent and use information in DHCP messages to populate its ARP cache. This mechanism, in combination with the DHCP address allocation process described above, can be
 520 advantageously utilized to control access to the network.

With reference to FIG. 7 (and FIG. 9), the CMTS acts as a DHCP relay agent and snoops on DHCP messages exchanged between the host acting as a DHCP client and the DHCP server. At step 707 in FIG. 7 and step 907 in FIG.

9, the DHCP server issues a DHCPACK message with configuration parameters including the committed IP address allocated to the subscriber to the particular service selected, as described above. This only occurs after the proper authentication steps have been performed. FIG. 11 is a flowchart illustrating the processing performed at the CMTS. At step 1101, the CMTS receives the DHCP message and proceeds to snoop on its contents in the process of relaying it to the proper DHCP client at step 1106. At step 1102, the CMTS determines that the DHCP message is a DHCPACK message by examining the options field of the DHCP message. At step 1103, the CMTS proceeds to read the information from the DHCPACK message, in particular the "yiaddr" and "chaddr" fields (fields 805 and 808 in FIG. 8) which indicate the assigned IP address and the hardware MAC address respectively. The CMTS updates the ARP cache to reflect the mapping of IP address to MAC address (and SID) used for communication with the subscriber. Thereafter, the CMTS permits upstream and downstream packets to and from the particular network access device based on the ARP cache entry. Where there has been a failure in authentication, this will result in a failure to update the ARP cache and in a denial of access to the access network infrastructure. The information in the ARP cache entry can also be utilized by the CMTS to monitor attempts to "spoof" the IP address of an authenticated subscriber: e.g., by refusing to forward upstream packets with a source IP address that is not associated with the proper SID or MAC address. The CMTS can set the timeout for the ARP cache entry to the same value as the timeout for the IP address specified in the DHCP exchange. If a customer subsequently releases the IP address by issuing a DHCPRELEASE message (or declines to accept the offered IP address, both at step 1104), the CMTS flushes the corresponding ARP cache entry upon receipt of the DHCP message at step 1105.

The DHCP server should be configured to ensure that every response that changes an IP address assignment or a lease on a DHCP assignment gets relayed to the CMTS. For example, a network access device using standard DHCP can send a DHCPRENEW message directly to the DHCP server. The DHCP server would respond directly to the client, and the CMTS would not

555 perform a DHCP relay function for these messages. Such direct communication between the DHCP client and server may modify information (e.g. lease time) used by the CMTS to manage the ARP cache. Therefore, it is advantageous to modify the DHCP server so that it will notify the CMTS of any changes to IP address assignment including lease time. This can be accomplished, for example, 560 by sending all responses to the CMTS to relay to the client—rather than sending them directly to the client.

It is advantageous for the access network infrastructure to support quality-of-support, in particular to differentiate service between subscribers to different services or service providers. For example, with reference to FIG. 1, a 565 subscriber to services offered by a service network 151 could get a higher share of access link capacity in the access network infrastructure than a subscriber to services offered by service network 152. Likewise, the aggregate capacity for subscribers to service network 151 may be different (e.g. higher) than aggregate capacity for subscribers to service network 152. It is preferable to avoid relying 570 on the network access devices 101...104 or the network interface units 111...114 to enforce service policies. Instead, it is advantageous for the access network infrastructure to have a policy engine, referred to herein as a policy decision point, and points at which policies may be enforced. The policy decision point, for example, could be a server in the service activation system 160 in the access 575 network infrastructure with access to the relevant registration information for subscribers. The policy decision preferably should be made at a point at which the association between the service network and the access network has been established, e.g. when the authenticated IP address is provided to the network access device as described above. The service class assignment would reflect the 580 results of negotiation between the entity operating the access network infrastructure and the entity operating the service network.

FIG. 12, for example, sets forth a timeline diagram of messages exchanged in assigning a service class to a subscriber, in the context of the particular HFC architecture described above. At 1201, the network access device 585 1210 sends the DHCPREQUEST message to the DHCP server 1250, requesting

the service-related IP address identified in a previous DHCP OFFER message. At 1202, the DHCP server 1250 commits to assigning the IP address to the network access device 1210 and transmits a DHCP ACK message containing the configuration parameters for the device. At 1203, the DHCP server 1250 sends a message to a Policy Decision Point ("PDP") 1240 authorizing assignment of a particular service class to the subscriber. The PDP 1240 transmits traffic parameters for the authorized service flow to the CMTS 1230, which receives the policy parameters at 1204 and begins the process of Dynamic Service Addition (DSA) in accordance with DOCSIS to schedule the service flow. A service flow, as defined by DOCSIS, is a unidirectional flow of packets that is provided a particular quality of service. At 1205, the CMTS 1230 sends a dynamic service addition request (DSA-REQ) message to the cable modem 1220 attached to the relevant network access device 1210. At 1206, the cable modem 1220 sends a dynamic service addition response (DSA-RSP) message after confirming that the cable modem 1220 can support the service flow. At 1207, the CMTS 1230 sends a dynamic service addition acknowledge (DSA-ACK) message after enabling transmission and reception of data on the new service flow. At 1208, the service flow has been allocated, and transmission on the new service flow has been enabled. At 1209, the network access device 1210 can begin transmitting data upstream to the CMTS 1230 which, as described above, forwards the data to the relevant packet-switched network.

Where the packet-switched network also provides for differentiation in service, e.g. based on the "DiffServ" framework, the cable modem 1220 can mark the Type Of Service (TOS) field of the packet to indicate the service class for the packet. See K. Nichols, S. Blake, F. Baker, D. Black, "Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers," IETF Network Working Group, RFC 2474 (December 1998); S. Blake, D. Black, M. Carlson, E. Davies, Z. Wang, W. Weiss, "An Architecture for Differentiated Services," IETF Network Working Group, RFC 2475 (December 1998), which are incorporated by reference herein. Using the policy information received from the PDP 1240, the CMTS 1230 can act as a policy enforcement

point and police the packets received from the cable modem 1220 and remark the TOS field of the packet where there is any deviation from the policy specified for the particular subscriber. Similarly, the CMTS 1230 can police service flows in the downstream direction, again based on the policy provided by the PDP 1240.

The CMTS 1230 performs scheduling based on the policies provided by the PDP 1240. Scheduling can be based on differentiation between subscribers to different services or service providers and can also be used to ensure that reserved minimum traffic rate requirements are met (this may require some measurement of capacity allocated to different subscribers as well as an accounting of bandwidth usage). One method of scheduling in the CMTS is to follow class based queuing (CBQ). See S. Floyd, V. Jacobson, "Link-Sharing and Resource Management Models for Packet Networks," IEEE/ACM Transactions on Networking, Vol. 3, No. 4, August 1995, which is incorporated by reference herein. CBQ allows for support of quality-of-support and flexible link sharing. The mechanism can be used to share capacity of a link across services, protocol families and/or traffic types. It can be used to allocate "shares" to individual service providers and subdivide the share amongst the subscribers of a given service provider. For example, arriving packet flows are aggregated into classes, each class having a "priority" and a throughput allocation. The traditional method is to use the information in the packet header, although it is possible, in the HFC architecture, to classify based on SID and the policy set up at the time the service flow is setup. A hierarchy of classes is constructed, e.g. as depicted in FIG. 13. The access link 1300 is shared between three service providers 1351, 1352, and 1353. Each service provider shares its allocation of link bandwidth among its customers, i.e., 1311, 1312, 1313 being subscribers of service provider 1351; 1321, 1322 being subscribers of 1352; 1331 being a subscriber of service provider 1353. The link scheduler may associate weights for each service provider. The service provider "weight" determines the overall share of access link bandwidth to all subscribers of the service provider, when the link is fully utilized. The link scheduler further allocates weights to each customer of the service provider—the customer's weight determining the share of the service provider's link bandwidth

available to the subscriber. It can then be possible to “borrow” bandwidth from other service provider subscribers when the service provider’s share is not fully
650 subscribed, as well as from other service providers when the link is not fully loaded. It is advantageous for the scheduling to provide such flexible sharing. When the link is underloaded, there is no need for any regulation of access to the HFC link by active SIDs. The scheduler can regulate a particular customer, e.g. by postponing giving grants to a SID of the customer, only when the customer is
655 over the limit in terms of bandwidth share during the short-term interval.

The foregoing Detailed Description is to be understood as being in every respect illustrative and exemplary, but not restrictive, and the scope of the invention disclosed herein is not to be determined from the Detailed Description, but rather from the claims as interpreted according to the full breadth permitted by
660 the patent laws. It is to be understood that the embodiments shown and described herein are only illustrative of the principles of the present invention and that various modifications may be implemented by those skilled in the art without departing from the scope and spirit of the invention. For example, the detailed description describes an embodiment of the invention with particular reference to
665 an HFC access network architecture. However, the principles of the present invention could be readily extended to other access network architectures, such as DSL, wireless, satellite, etc. Such an extension could be readily implemented by one of ordinary skill in the art given the above disclosure.